Opinion Polarization by Learning from Social Feedback[1]

S. Banisch and E. Olbrich
Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

arXiv:1704.02890v1 [physics.soc-ph] 7 Apr 2017

Abstract

We explore a new mechanism to explain polarization phenomena in opinion dynamics. The model is based on the idea that agents evaluate alternative views on the basis of the social feedback obtained on expressing them. A high support of the favored and therefore expressed opinion in the social environment, is treated as a positive social feedback which reinforces the value associated to this opinion. In this paper we concentrate on the model with dyadic communication and encounter probabilities defined by an unweighted, time-homogeneous network. The model captures polarization dynamics more plausibly compared to bounded confidence opinion models and avoids extensive opinion flipping usually present in binary opinion dynamics. We perform systematic simulation experiments to understand the role of network connectivity for the emergence of polarization.

1 Introduction

The public discourse around political polarization has regained momentum in the recent past. The fundamental divide between Democrats and Republicans in the United States and even more so the heatedness of some of the confrontations are often expounded with reference to an increase in differences in political attitudes that, seemingly, has taken place during the last decade. The rise of nationalistic voices in the European sphere, likewise, can be seen as a reflection of fundamental differences in attitudes towards the European project and different ideas regarding the »right« direction to go. These developments are very challenging from the socio-political point of view, and – given the complex interplay of a great variety of actors and entities at different societal levels – they are also a hard puzzle for the social sciences. Many different pieces must be brought into relation in order to come to an encompassing scientific understanding: from empirical measurements of phenomena that we intuitively relate to the problem to an analytical exploration and filtering of possible causal mechanisms that explain their coming-about.

With this paper, we aim at contributing to a recent development in analytical sociology to identify mechanisms and conditions for the emergence of a polarization. In this context, polarization refers, one the one hand, to a process by which such strong divergences of political opinions come about, on the other, it refers to a measure of the extend to which a certain distribution of opinions is considered as polarized. The persistence and emergence of a polarized situation still poses a challenge in theoretical sociology and seeking microscopic explanations by building mathematical or computational models of how individuals act and exchange opinions has a long tradition in the field (French Jr, 1956; Abelson, 1964; DeGroot, 1974; Chatterjee and Seneta, 1977; Friedkin and Johnsen, 1990; Deffuant et al., 2000; Hegselmann et al., 2002). In this paper, we provide a parsimonious answer to Abelson's old question of »what on earth one must assume in order to generate the bimodal outcome of community cleavage studies« (Abelson, 1964, p. 153).

The main idea is that individuals express their opinion about an issue and are sensitive to approval and disapproval by their peers. Agreement leads to a positive experience which strengthens support on the expressed opinion. Conversely, disagreement is assumed to be related to a negative sentiment and decreases support of the current opinion. In a wide range of interaction networks this reinforcement leads to the formation of opinion clusters in which agents become collectively more and more convinced of the respective group opinion. The process we propose

[1]This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 732942 (www.Odycceus.eu).

can be seen as an abstraction from recent models of polarization that rely on ideas from argument persuasion (Dandekar et al., 2013; Mäs and Flache, 2013). However, our model comes with another connotation as the opinion change process is not assumed to involve argument processing but change is mediated by the positive (negative) experience that agreement (disagreement) brings about.

2 Modeling Opinion Polarization

Modeling opinion polarization is recently receiving considerable attention in the opinion dynamics community. In the last few years, a series of possible interaction mechanisms have been proposed by which a population polarizes and approaches a state with a bi-polar opinion distribution (Macy et al., 2003; Baldassarri and Bearman, 2007; Flache and Macy, 2011; Dandekar et al., 2013; Mäs and Flache, 2013; Mäs and Bischofberger, 2015) and researchers have started to assess the relevance of these competing explanatory variants through experiments (Moussaïd et al., 2013; Chacoma and Zanette, 2015; Takács et al., 2016). The comparison of different models in terms of experimentally-founded microscopic assumptions on the one hand and their capabilities to generate plausible macroscopic outcomes on the other is actually a quite remarkable scientific program as it aims at an empirical justification of rather different links in the explanatory chain (Hedström and Ylikoski, 2010).

A major motivation for such a mechanism-based approach to polarization has been that fact that most early models of opinion dynamics lead to »universal ultimate agreement« (Abelson, 1964, p. 153). These models are generally based on mechanisms of *positive social influence* by which people's views on an issue tend to become more similar in interaction[2]. Without further ingredients, opinion models in which agents assimilate their views in interaction predict that a population will converge to a state of complete consensus (cf. Dandekar et al. (2013)).

The second generation of models of opinion dynamics try to address this by including further constraints in the models, most prominently, assumptions about value or *opinion homophily* (McPherson et al., 2001). In the opinion dynamics context this type of homophily refers to the fact that similar views may lead to attraction and an increasing likeliness of interaction. Paired with positive social influence it describes a process of repeated influence events by which »similarity leads to interaction and interaction leads to still more similarity« (Banisch et al., 2010, 109). In the classical models with bounded confidence (Hegselmann et al., 2002; Deffuant et al., 2000) this is implemented via a threshold mechanism that switches off influence if the opinion discrepancy exceeds a certain threshold. Similar mechanisms have been proposed for models with an vectorial opinion representation (e.g. Axelrod (1997); Banisch et al. (2010)).

Bounded confidence can lead to a stable co-existence of a plurality of opinions in a population. However, unless »extreme« agents are artificially inserted (Deffuant et al., 2002), bounded confidence models cannot drive a population into the extremes or lead to the emergence of antagonistic opinion groups as the opinion averaging procedure always locally reduces diversity. It crucially depends on the initial opinion diversity whether or not multiple opinions survive and the model will always end up with final opinions that are more moderate than the initial extremes.

The recent approaches to modeling polarization dynamics introduce new mechanisms of opinion exchange and influence in addition to positive influence and homophily. They can be roughly classified into models that rely on the assumption of *negative social influence* and models that draw upon ideas from *persuasion theory* (cf. Mäs and Bischofberger (2015)).

The former seek an explanation of polarization patterns by assuming a negative social influence in the interaction of distant agents such that the encounter and communication of two agents with very different views leads them to adopt even more distant positions in the opinion space (Mark,

[2]See (Takács et al., 2016) for a recent experimental confirmation of opinion assimilation

2003; Macy et al., 2003; Baldassarri and Bearman, 2007; Flache and Macy, 2011). Despite the possibility to generate stable polarization patterns driving the population into two maximally extreme clusters, a recent experimental study found no indication for negative influence of this kind (Takács et al., 2016).[3] It is noteworthy that the combination of positive and negative social influence can lead to polarization but that this feature disappears in the presence of opinion homophily (Mäs and Bischofberger, 2015).

The second type of models capable of explaining polarization is based on psychological work on persuasion and attitude change (Fishbein, 1963; Lord et al., 1979; Petty et al., 1981; Ajzen, 2001). Persuasion models generally assume that communication partners exchange arguments about the object on which an attitude is formed and that new arguments are learned from an interaction partner if they are in support of an agents view. The model by Mäs and Flache (2013) is based on an explicit representation of attitudes borrowed from value-expectancy theory which treats an attitude as a weighted combination of a set of pro- and con-arguments regarding certain aspects of the issue under discussion (see also Urbig and Malitz (2005) for an early related model). Their so-called »argument-communication theory of bi-polarization« posits a mechanism for the co-evolution of arguments and the associated weights with the latter encoding a limited capacity to process all arguments at once. That is, only a subset of arguments is considered to be relevant and recently discussed issues play a more important role. At each time step, two agents are chosen at random with a probability that is proportional to the similarity of their attitudes (opinion homophily). In the interaction, the first agent adopts a randomly chosen argument from its interaction partner (positive social influence). The larger homophily is, the larger is the likeliness that the adopted argument speaks in favor of their own attitude, and if the similarity in attitudes of both agents came about by different argument subsets (due to the recency weightings) it becomes likely that a new supportive argument is learned. As a result, the model leads to a process where agents with a similar attitude mutually reinforce that attitude by the exchange of supportive arguments.

Another proposal to the modeling of bi-polarization has been made by Dandekar et al. (2013) who base their argument on the work of Lord et al. (1979) on *biased assimilation*. The principle idea is that people who are very convinced of their view tend »to accept "confirming" evidence at face value while subjecting "disconfirming" evidence to critical evaluation«. This is sometimes also referred to as confirmation bias. Unlike the model described above, Dandekar et al. (2013) take an abstract perspective in terms of opinion representation and include biased argument evaluation into the classic repeated averaging model (DeGroot, 1974) which operates on a single continuous opinion dimension ($o_i \in [0, 1]$). A bias function is introduced that decides about strength and direction of opinion change as a function of the current conviction of one or the other view (the extremes 0 and 1 are interpreted as two opposing opinions on an issue). The model contains DeGroot's averaging process as a limiting case and Dandekar et al. (2013) shows that there is a critical bias level at which the process becomes polarizing. Noteworthy, polarization may occur even in the absence of opinion homophily.

This paper adds a new mechanism to this literature. Like in persuasion models, opinions are reinforced by encountering agents with similar views, but in our model this »reinforcement in agreement« is mediated by a very different psychological process. Namely, reinforcement or respectively a weakening of support of their expressed view is not obtained by a costly process of argument persuasion but rather by the positive (negative) experience that agreement (disagreement) brings about. Agents form their opinions on the basis of the social feedback they obtain by expressing them.

[3] The experiment was indicative of a negative influence only when people held very similar opinions, which can be seen as a tendency to individualization. Assumptions in line with this finding have been investigated in, for instance, Banisch (2010), and generally lead to a final opinion distribution characterized by small fluctuations around the group consensus points.

3 Model Description

3.1 Theoretical Model

Suppose there are two opinions that agents can adopt and express. We denote these alternative options as o_i, i being an agent index, and set $o_i \in \{-1, 1\}$ for further convenience. In the opinion model we put forth here, an agent (say i) is chosen at random and expresses his current opinion o_i to a randomly chosen neighbor j. This agent responds on i's expression with her current opinion o_j which can signal either agreement or disagreement in the two-opinion case.
We further assume that j's agreement resonates positively her disagreement negatively in i. This is accounted for by two real-valued terms per agent – $Q_i(1)$ and $Q_i(-1)$ – that capture how well the expression of 1 and -1 respectively is perceived by peers in an agent's social environment. That is, the $Q_i(o)$ represent an internal evaluation of the different options based on the social response the agent obtains on expressing them. These values are updated as

$$Q_i(o) \leftarrow \begin{cases} (1-\alpha)Q_i(o) + \alpha r_i : \text{ if } o = expression \\ Q_i(o) : \text{ else.} \end{cases} \tag{1}$$

with

$$r_i = o_i o_j \tag{2}$$

leading to a positive feedback for $o_i = o_j$ and to a negative one if $o_i \neq o_j$. The parameter α is referred to as learning rate (see below) and governs the magnitude of change. It is set to $\alpha = 0.05$ for the simulations performed in this paper.
On expressing their current opinion agents thus receive affirmative or non-confirming response depending on the current opinion in their neighborhood. Agreement leads to a positive experience ($r_i = 1$) by which the evaluation $Q_i(o)$ of the respective expressed opinion increases. Conversely, a disconfirming response decreases the value associated to the respective opinion. We shall therefore interpret the values $Q_i(o)$ as associated to the two opinions as a strength with which the respective view is supported by i or likewise as i's conviction regarding the two alternatives. Consequently, we assume that when asked to articulate his current opinion or to respond to such an articulation, agents choose to express that option which (s)he more strongly supports at the current time step. In other words, $o_i = -1$ if $Q_i(-1) > Q_i(1)$ and $o_i = 1$ if $Q_i(-1) < Q_i(1)$ or, more generally:

$$o_i = \arg\max_o Q_i(o). \tag{3}$$

However, we assume that agents generally tend to follow (3) in their expression choice, but deviate from this scheme with a small probability $\epsilon = 0.1$. The mathematical reason for this is that we make sure, in this way, that both opinion options are tested by the agents so that both Q-values are actualized from time to time; for this reason ϵ is usually referred to as exploration rate. On the other hand, exploration seems also plausible especially when the difference in conviction is small[4]. Moreover, one may argue that the noise introduced by ϵ accounts for certain issues that might occur in communication such as misunderstanding or misinterpretation of an articulation.
Readers familiar with game theory and reinforcement learning will have realized that this model set-up casts opinion dynamics as a game played repeatedly on a network in which agents learn the best response by a simple form of independent Q-learning (Sutton and Barto (1998); Busoniu et al. (2008)). In this context, opinion expression should be seen as an action that leads to a certain reward r and the (state-less) Q-values associated to the two possible actions (i.e. $o_i \in \{-1, 1\}$) are updated based on this reward signal. In fact, we will make use of this analogy

[4]In fact, this would motivate the so-called softmax action selection (Sutton and Barto, 1998) which assigns equal probability to the two options when there Q values are equal and gradually favors the option with larger Q the more they differ.

in the sequel to mathematically characterize a series of stylized situations that will be helpful to provide an overall picture of the model behavior.
We also note that the model behavior is essentially determined by the way in which the reward system is designed. In this basic model agents learn to avoid dissonance and play a coordination game, but we envision that different more complex and possibly heterogeneous rewards are a propelling ingredient that deserves further exploration. The main purpose of this paper, however, is to introduce this type of opinion game in the context of opinion models and polarization dynamics in particular, and to demonstrate that opinion polarization can result from very few relatively mild basic assumptions:

1. agreement is positively, disagreement negatively experienced,
2. these experiences drive opinion conviction (ΔQ_i) and expression (o_i),
3. the probability of interaction is structured.

While the first two assumptions are realized through rewards and the update scheme as specified by (1) - (3), the third aspect is included in the model by an interaction network that determines the probability with which pairs of agents engage in communication with one another. For this purpose, we generate a random geometric network (Penrose, 2003; Dall and Christensen, 2002) to define the neighborhood structure. According to that network model, N agents are assigned a random position in the unit plane $(x_i, y_i) \in [0,1] \times [0,1]$ and a link is established whenever the distance between two agents is below a threshold r. Notice that the density of the network depends on both the number of agents and the threshold r and can be approximated as $\pi N r^2$ for large N (Penrose, 2003). Notice also that random geometric graphs are a special type of random graphs that differ from the latter especially regarding clustering (Dall and Christensen, 2002).

3.2 Implementation Details

```
r = 0.175;  % distance threshold for random geometric graph
[linkList,pos] = spatialRandomGraph(N,r); % generate spatial random graph

alpha = 0.05; % learning rate
eps = 0.1; % exploration rate

Q = rand(N,2)-0.5; % random initial values in [-0.5,0.5] for both opinions

for step = 1:steps % iterate as long as indicated by steps

    a1 = randi(N);  % random choice of agent that expresses his opinion
    expression = (Q(a1,2) > Q(a1,1));   % expression (1 -> o = 1 | 0 -> o = -1)
    if(rand() < eps)
        expression = -(expression - 1); % exploration
    end;

    numN = linkList(a1,1);              % random choice of a responding agent
    a2 = linkList(a1,randi(numN)+1);    % out of the neighbors of a1
    reaction = (Q(a2,2) > Q(a2,1));     % response opinion articulated by a2

    reward = (expression*2-1) * (reaction*2-1); % agreement -> 1 | disagreement -> -1

    % update of Q-values for a1 considering what has been expressed
    Q(a1,expression+1) = (1-alpha)*Q(a1,expression+1) + alpha * reward;

end;
```

Figure 1: Runable MatLab-Code of the model.

In order to facilitate the reproduction of results and give readers the possibility to make their own explorations with the model, we provide all the details of the model in form of the two code snippets shown in Figure 1 and 2. Noteworthy, the model is simple enough to be implemented in

```
function [linkList,pos] = spatialRandomGraph(N,r)

pos     = rand(N,2); % random positions in [0,1]x[0,1]
linkList = [];       % empty list for neighbors

% go through all possible pairs and check if linked
for a1 = 1:N
    linkList(a1,1) = 0;
    for a2 = 1:N
        % euklidean distance smaller than r (and a1 != a2)?
        dist = sqrt((pos(a1,1)-pos(a2,1))^2 + (pos(a1,2)-pos(a2,2))^2);
        if(ne(a1,a2)&&dist < r)
            % count number of neighbors in first element of linkList
            linkList(a1,1) = 1 + linkList(a1,1);
            % add neighbor
            linkList(a1,linkList(a1,1)+1)=a2;
        end;
    end;
end;

end
```

Figure 2: Runable MatLab-Code of the random graph model.

a few lines in MatLab (R2013a for Mac) and is executable as it is presented (N and the number of iterations `steps` have to be specified). Figure 1 presents the model initialization and the iteration loop, Figure 2 the random geometric graph model. Notice that for the computation-intensive experiments the model was also implemented in C++.
In addition, two online implementations of the model are available under

- `www.universecity.de/demos/OpinionValuesSmall.html`
- `www.universecity.de/demos/OpinionValuesBig.html`

The former is a model with $N = 100$ agents which is also described in the next section. The latter is an implementation with $N = 10000$ agents where the network is not visualized for performance reasons. Notice that the models only run in Browsers with WebGL support.

4 Model Behavior on Random Geometric Graphs

The aim of this section is to provide an intuition about the model behavior. We will look in some detail on realizations of the model on random geometric graphs. We will consider individual trajectories as well as some macroscopic indicators and show that polarization patterns emerge in a wide range of networks.

4.1 Simulation Setting

As described above we generate a random spatial network to define the neighborhood structure of the agents. N agents are assigned random position in the unit plane $(x_i, y_i) \in [0, 1] \times [0, 1]$ and a link is established whenever the distance between two agents is below a threshold r. The agent population is initialized by setting the initial values $Q_i(1), Q_i(-1)$ at random according to a uniform distribution within $[0, 1]$. This defines the initial opinions (i.e. what the agents express at $t = 0$). Due to the initialization of the Q–values, on average, half of the population will initially hold opinion 1 and the other half opinion -1.
The exemplary model realization discussed in this section is done with $N = 100$ agents. The threshold for the random geometric graph model is $r = 0.175$ and the resulting network is shown in Fig. 3. The learning rate is $\alpha = 0.05$ and the exploration rate $\epsilon = 0.1$. Both values are used throughout all the analyses presented in this paper.

4.2 Macroscopic and Microscopic Dynamics

An example run of the model with $N = 100$ agents is shown in Fig. 3. It shows for 10000 simulation steps the microscopic system configuration at times zero, 2500, 5000, 7500 and 10000 along with different macroscopic observables that measure the amount of polarization in a population. We follow DiMaggio et al. (1996) in the definition of polarization measures and define *dispersion* as the variance over the distribution of convictions $\Delta Q_i = Q_i(1) - Q_i(-1)$ and *bimodality* by the kurtosis rescaled to the interval $[0, 1]$. In addition, the polarization measure introduced in Flache and Macy (2011) referred to as *dissimilarity* is shown. Moreover, the *average opinion* (fraction of agents that express $o_i = 1$) and the *average strength of support* within the two different groups of supporters is shown over time.

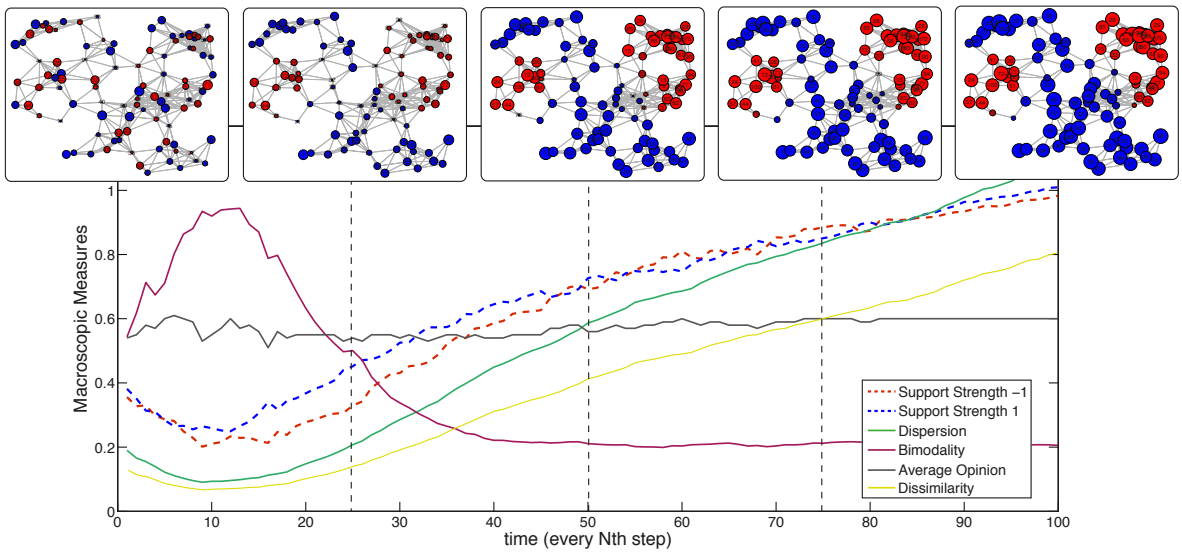

Figure 3: Time evolution of the model. The plot shows 10000 iterations for 100 agents (on average 100 expressions by each agent) along with the system state at $t = 0, 2500, 5000, 7500, 10000$. The average opinion represents the fraction of agents that support 1 and the support strength is the value difference averaged over the respective sets of supporters. In the definition of the polarization dispersion and bimodality measures we follow DiMaggio et al. (1996).

Initially, due to the random initialization of values, approximately one half of the population supports *expression* 1 and -1 is supported by the other half. Notice that these fractions do not change a lot during the simulation even if single agents do change their opinion. The left-most network shows that supporters are distributed without a particular spatial organization. Average strength of support for both alternatives is at the same level and relatively low.
The strength of support is decreasing during an initial period of alignment which results in a strong spatial organization into opinion regions is clearly visible at step 2500. This spatial distribution of opinions remains rather stable in the subsequent steps. While strength of support is still small for most agents at $t = 2500$ (approximately at the initial level), some clusters have emerged in which one opinion is more strongly supported after 5000 steps. In particular, we observe two regions in which agents strongly support 1 (blue) which are connected by a bend of agents that moderately support 1. We also observe two disconnected clusters where -1 (red) is supported more and more strongly. However, decisive changes may still occur at the interfaces between regions in which different opinions are supported and some clusters may be invaded in a long transient. Despite support strength remaining small in those areas, on average, support strength increases for both opinions.
This evolution is also captured by the different polarization measures. First, dispersion and

dissimilarity behave very similarly and show an initial decrease in polarization followed by a steady increase once the opinion clusters have formed. Kurtosis – capturing the bimodality of a distribution – behaves differently. First, one should notice that large kurtosis values are indicative of low bimodality. The first peak in the measure is therefore indicative of a considerable reduction in polarization. However, as with the other measures, it indicates that polarization increases to the initial level after approximately 2500 steps (second network). It further decreases but, opposed to the other measures, the bimodality indicator reaches a stable level at around 4000 to 5000 steps. This is due to the fact that kurtosis measures only whether a distribution is bimodal but does not account for the distance between the two peaks of the distribution.
To fully understand the dynamical behavior of the model, let us look at the temporal evolution of some particular agents. In Fig. 4 the values $Q_i(1)$ (blue) and $Q_i(-1)$ (red) are shown for three different agents. Along with the values, the events where the ranking of values changes so that another opinion will be chosen in expression and response are highlighted by yellow stars. The first one (agent 29, located in the upper right corner of the networks in Fig. 3) starts in a state where $Q(1) > Q(-1)$ but both values are at a relatively high level. Initially, this agent is surrounded by an equal number of blue and red agents which leads to an decrease in both $Q(1)$ and $Q(-1)$ as the expected reward (local feedback) is zero in such a situation. However, already at step 2500 (see second network in Fig. 3) the neighborhood of 29 has aligned to -1 (red) and from that time on the value $Q(-1)$ increases steadily until it saturates at $Q(-1) = 1$ which is the expected social reward in an homogeneous surrounding. Moreover, $Q(1)$ decreases step by step as the result of negative feedback on expressing 1 due to the exploration probability ($\epsilon = 0.1$). $Q(1)$ will eventually approach -1.
This behavior is characteristic for many agents in the simulation without a too strong initial preference. There is a initial phase of alignment in which a local consensus emerges – weakly supported, at first – which, once established, leads to an reinforcement of the respective view. Notice that agents with a strong or moderate initial support of red in that cluster do in fact never change opinion. Their values both tend to zero as well but the positive reinforcement of their initial preference begins before a change of ranking takes place.

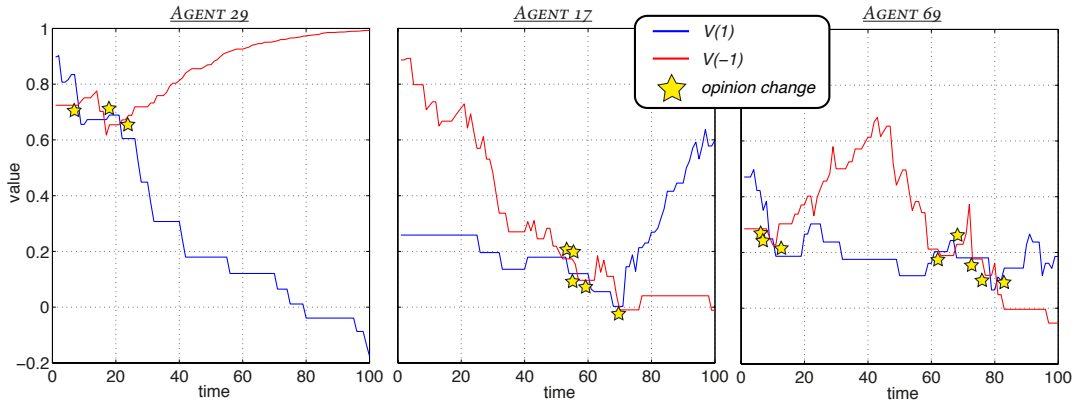

Figure 4: Time evolution of three selected individuals in the above example run. The plot shows 10000 (on average 100 expressions by each agent). The yellow stars represent opinion change events where the values associated to the two alternatives change ranking.

Secondly, we consider agent 17 with a relatively strong initial preference for -1. Her neighborhood remains unaligned for a rather long time and still is at step 5000. In other words, the agent is located in the interface region between different opinion clusters that have emerged early in the simulation. As shown in Fig. 3, the agent forms part of a small cluster (along with 62, 75, 76) that is only slowly »invaded« by the blue opinion. Around step 7000, however, most of her neighbors express preference for opinion 1 which leads to a rapid increase of $Q(1)$ in subsequent steps.

Finally, another node at the interface between two opinion regions is considered, namely agent 69. With an initial preference for blue, this agent first aligns with the red cluster to which agent 29 belongs and learns to support this opinion rather strongly. However, as the neighboring cluster around agent 17 aligns more and more on 1, expressing -1 yields a negative feedback with higher probability and leads to an decrease of $Q(-1)$. Consequently, in between step 6000 to 8000 agent 69 experiences a second period of opinion changes the result of which is a slight preference of 1. It is likely that a weak to moderate preference for blue will be stable as the agent is surrounded by 4 red and six blue agents. However, there may be chains of random events that take the agent to temporarily supporting -1 again.

An important characteristic visible in this microscopic perspective is that agents do not immediately adapt their expressed opinion as a response to interaction with unaligned peers. The imitation process usually implemented in binary opinion dynamics is replaced by a continuous re-evaluation of the two options. Periods of »indecision« during which individuals try different options alternate with periods during which agents show a clear preference of one or the other opinion.

Yet, this continuous adaptation is different from continuous models for opinion dynamics where an »average opinion« emerges within different network clusters. To the contrary, the social feedback mechanism produces polarization by reinforcement of the support assigned to an opinion once a community is locally aligned. In other words, the contraction mechanism leading to different opinion clusters in bounded confidence models (at least if the initial diversity is large enough) is replaced by a mechanism that drives opinion clusters to a higher degree of polarization (even if initial opinions are relatively close).

4.3 Network Density and Consensus Probability

We now characterize the model on random geometric graphs by looking at its long-run behavior as a function of the network density. Notice again that the network density relates to the threshold r approximately as $\pi N r^2$ (Penrose, 2003) and density is in fact more suited to compare systems of different sizes (N). We consider the consensus probability (the fraction of realizations that end up with all agents in the same opinion state) and the dissimilarity polarization index (Flache and Macy, 2011) to identify the range of networks in which polarization may be observed. Fig. 5 shows this for three different system sizes. Each data point is obtained by averaging over 100 realizations.

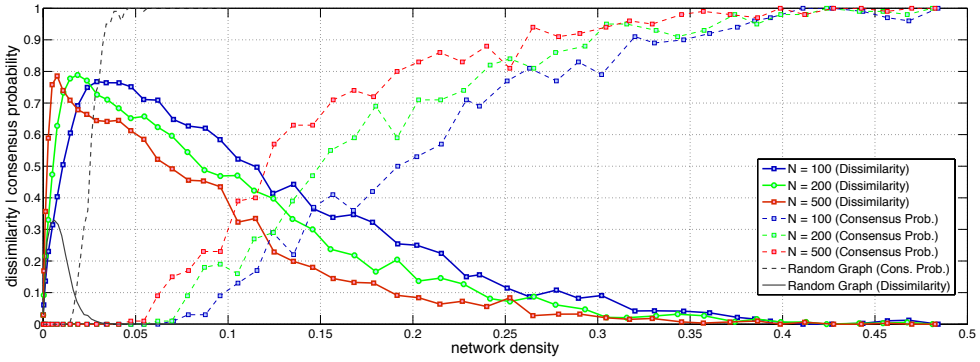
dissimilarity I consensus probability
network density
N = 100 (Dissimilarity)
N = 200 (Dissimilarity)
N = 500 (Dissimilarity)
N = 100 (Consensus Prob.)
N = 200 (Consensus Prob.)
N = 500 (Consensus Prob.)
Random Graph (Cons. Prob.)
Random Graph (Dissimilarity)

Figure 5: Probability of consensus (all agents maximally support the same opinion) and polarization (measured as dissimilarity) as a function of the density of the random geometric graph.

The figure shows that pronounced polarization patterns are very likely in a considerable range of density values. The low polarization as $\langle k \rangle \to 0$ is explained by the fact that it leads to a high probability of isolated nodes which escape the feedback process and stay at their initial

support level. As density increases, the likeliness of components in which the social feedback process plays out quickly increases leading to a high probability and value of polarization. As the consensus probability (all agents maximally support the same opinion) is zero up to $\langle k \rangle \approx 0.05$, all realizations led to the formation of distinct opinion groups with extreme support for their view. As density further increases the polarization indicator decreases due to the fact that a proportion of realization approach a consensus state. The example considered in the previous section falls into that regime. We also consider in Fig. 5 the convergence behavior on a random graph without spatial organization (gray lines) and observe that the transition to full consensus is rather sharp in this case ($N = 200$). This indicates that the spatial organization into relatively segregated communities of agents (see the networks in Fig. 3) is decisive for polarization to emerge. This interpretation is supported by game-theoretic considerations and points at the community structure of a network as the most important factor to explain polarization in the context of our model (both issues are briefly discussed in the next section).

5 Remarks

Mathematical Characterizations. The proposed opinion model is strongly rooted in reinforcement learning and can be seen as a repeated stochastic game played on a network. We treat opinion expression as action and the social feedback as the payoff that an expression receives. We want to highlight that there is actually a rich body of game-theoretic literature which studies such kind of games on networks (see Jackson and Zenou (2014) and references therein) and coordination games – as specified by our reward function (2) – have received particular attention. For instance, it is easy to show by game-theoretic arguments that two weakly connected groups maintain opinion diversity (gate keeping at structural holes). While the application of these results to the model is promising and will be presented if our contribution is accepted, they go beyond the constraints of this submission.

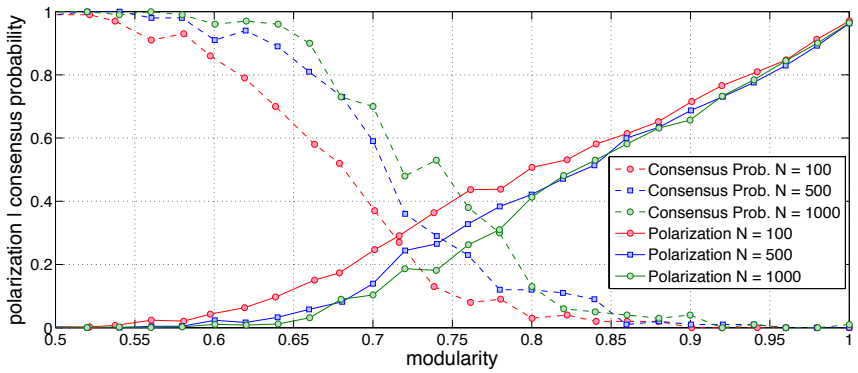

Figure 6: Probability of consensus and polarization (measured as dissimilarity) as a function of network modularity.

Community Structure. The model presented here defines a polarization process on a fixed network and allows to extract socio-structural conditions for the emergence of polarization. Further analyses with the so–called island graph model have revealed that community structure of the network is the most important structural factor. Island graphs are constructed by defining a set of communities and varying the probability of within- and across-community connections. An analysis for graphs consisting of 10 equally sized communities is shown in Fig. 6 These results indicate an almost linear relationship between modularity and polarization. The game-theoretic analyses mentioned above are largely concerned with the identification of stable cohesive sets which are closely related to the concept of communities in networks. We are currently exploring

this link. In addition, we envision that the proposed model is well-suited as a tool for community detection and the identification of bridges (structural holes) in networks.

6 Conclusion

This paper introduces a model for opinion polarization that is based on a minimal set of assumptions. On the one hand, the proposed mechanism captures a reinforcement process that is similar to recently proposed persuasion models in that opinion–confirming interactions reinforce commitment to that view whereas disconfirming interaction weakens opinion support. On the other hand, the model is based on a reinforcement procedure which does not assume any more complex kind of argument processing on the side of the agents but instead conceives opinion evaluation and consolidation as the result of positive and negative experiences that are assumingly related to agreement and disagreement with peers. The model should be seen as a »toy model« in that regard and implications for real opinion processes be drawn with careful attention to its abstractness.

The proposed mechanism is meant to consider one potential source of opinion change in isolation: that is, the formation of strong convictions in an opinion based on how acceptable it is to advocate that standpoint in a social neighborhood. Rejection and confirmation of expressed opinions by peers lead to a re-evaluation of the opinion options which is mediated via a reward signal. It is noteworthy, that a bunch of recent studies in neurobiology suggest that »social influence mediates very basic value signals in known reinforcement circuity« (Campbell-Meiklejohn et al., 2010). While neurobiology posits that the »rewarding properties of social behavior may have evolved to facilitate group cohesion and cooperation« (Ruff and Fehr, 2014), our model suggests that human ability to coordinate with in-groups comes at the expense of a likely alienation to out-groups. Empirical and experimental research is needed to assess the relevance of these mechanisms for opinion dynamics.

In conclusion, we believe that linking opinion dynamics with social feedback mechanisms bears a great potential for modeling opinion formation processes in different media settings and the model presented here should be understood as an initial implementation of this paradigm.

Abelson, R. P. (1964). Mathematical models of the distribution of attitudes under controversy. *Contributions to mathematical psychology*, 14:1–160.

Ajzen, I. (2001). Nature and operation of attitudes. *Annual review of psychology*, 52(1):27–58.

Axelrod, R. (1997). The dissemination of culture: A model with local convergence and global polarization. *The Journal of Conflict Resolution*, 41(2):203–226.

Baldassarri, D. and Bearman, P. (2007). Dynamics of political polarization. *American sociological review*, 72(5):784–811.

Banisch, S. (2010). Unfreezing Social Dynamics: Synchronous Update and Dissimilation. In Ernst, A. and Kuhn, S., editors, *Proceedings of the 3rd World Congress on Social Simulation (WCSS2010)*.

Banisch, S., Araujo, T., and a, J. L. (2010). Opinion dynamics and communication networks. *Advances in Complex Systems*, 13(1):95–111. ePrint: arxiv.org/abs/0904.2956.

Busoniu, L., Babuska, R., and De Schutter, B. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, And Cybernetics-Part C: Applications and Reviews, 38 (2), 2008*.

Campbell-Meiklejohn, D. K., Bach, D. R., Roepstorff, A., Dolan, R. J., and Frith, C. D. (2010). How the opinion of others affects our valuation of objects. *Current Biology*, 20(13):1165–1170.

Chacoma, A. and Zanette, D. H. (2015). Opinion formation by social influence: from experiments to modeling. *PloS one*, 10(10):e0140406.

Chatterjee, S. and Seneta, E. (1977). Towards consensus: some convergence theorems on repeated averaging. *Journal of Applied Probability*, pages 89–97.

Dall, J. and Christensen, M. (2002). Random geometric graphs. *Physical Review E*, 66(1):016121.

Dandekar, P., Goel, A., and Lee, D. T. (2013). Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences*, 110(15):5791–5796.

Deffuant, G., Amblard, F., Weisbuch, G., and Faure, T. (2002). How can extremism prevail? a study based on the relative agreement interaction model. *Journal of artificial societies and social simulation*, 5(4).

Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3(01n04):87–98.

DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121.

DiMaggio, P., Evans, J., and Bryson, B. (1996). Have american's social attitudes become more polarized? *American Journal of Sociology*, 102(3):690–755.

Fishbein, M. (1963). An investigation of the relationship between beliefs about an object and the attitude toward that object. *Human relations*.

Flache, A. and Macy, M. W. (2011). Small worlds and cultural polarization. *The Journal of Mathematical Sociology*, 35(1-3):146–176.

French Jr, J. R. (1956). A formal theory of social power. *Psychological review*, 63(3):181.

Friedkin, N. E. and Johnsen, E. C. (1990). Social influence and opinions. *Journal of Mathematical Sociology*, 15(3-4):193–206.

Hedström, P. and Ylikoski, P. (2010). Causal mechanisms in the social sciences. *Annual Review of Sociology*, 36:49–67.

Hegselmann, R., Krause, U., et al. (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3).

Jackson, M. O. and Zenou, Y. (2014). Games on networks. *Handbook of Game Theory*, page 95.

Lord, C. G., Ross, L., and Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of personality and social psychology*, 37(11):2098.

Macy, M. W., Kitts, J. A., Flache, A., and Benard, S. (2003). Polarization in dynamic networks: A hopfield model of emergent structure. *Dynamic social network modeling and analysis*, pages 162–173.

Mark, N. P. (2003). Culture and competition: Homophily and distancing explanations for cultural niches. *American sociological review*, 68(3):319–345.

Mäs, M. and Bischofberger, L. (2015). Will the personalization of online social networks foster opinion polarization? *Available at SSRN 2553436*.

Mäs, M. and Flache, A. (2013). Differentiation without distancing. explaining bi-polarization of opinions without negative influence. *PloS one*, 8(11):e74516.

McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444.

Moussaïd, M., Kämmer, J. E., Analytis, P. P., and Neth, H. (2013). Social influence and the collective dynamics of opinion formation. *PLoS One*, 8(11):e78433.

Penrose, M. (2003). *Random geometric graphs*. Number 5. Oxford University Press.

Petty, R. E., Cacioppo, J. T., and Goldman, R. (1981). Personal involvement as a determinant of argument-based persuasion. *Journal of personality and social psychology*, 41(5):847.

Ruff, C. C. and Fehr, E. (2014). The neurobiology of rewards and values in social decision making. *Nature Reviews Neuroscience*, 15(8):549–562.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.

Takács, K., Flache, A., and Mäs, M. (2016). Discrepancy and disliking do not induce negative opinion shifts. *PloS one*, 11(6):e0157948.

Urbig, D. and Malitz, R. (2005). Dynamics of structured attitudes and opinions. In *Third Conference of the European Social Simulation Association*, pages 5–8. Citeseer.